

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

Claim 1 (previously presented): A computer-implemented method of collecting statistics in a database management system comprising a plurality of collection objects having a plurality of tables, at least one particular collection object of the plurality of collection objects comprising a plurality of particular tables of the plurality of tables, comprising:

receiving a workload comprising a plurality of database queries referencing the plurality of tables, at least two of the database queries that are different referencing a same table of the plurality of tables;

for each database query of the plurality of database queries of the workload:

identifying query statistics for said each database query of the plurality of database queries, the query statistics comprising a number of references to at least one table of said each database query and a number of occurrences of said each database query; and

consolidating, for said each database query, the query statistics to provide consolidated statistics comprising at least one cumulative table score for said at least one table, respectively, wherein said at least one cumulative table score for said at least one table is based on a previous at least one cumulative table score, if any, for said at least one table and a product of the number of references to said at least one table of said each database query and the number of occurrences of said each database query; and

generating a plurality of statistics collection tasks to collect database statistics to be used by an optimizer based on the consolidated statistics, the plurality of statistics collection tasks to collect the database statistics on the plurality of tables referenced by the plurality of database queries of the workload, the plurality of statistics collection tasks being based on a plurality of collection objects comprising the plurality of tables referenced by the plurality of database queries of the workload, a plurality of collection object scores are determined for the

ok to enter

plurality of collection objects comprising the plurality of tables referenced by the plurality of database queries of the workload, each collection object score for each collection object is based on said at least one cumulative table score for said at least one table of said each collection object, wherein the plurality of statistics collection tasks are ordered based on the plurality of collection object scores.

Claim 2 (canceled).

Claim 3 (canceled).

Claim 4 (previously presented): The method of claim 1 wherein the collection objects are table spaces.

Claim 5 (canceled).

Claim 6 (canceled):

Claim 7 (currently amended): The method of claim 1 wherein the [[a]] at least one cumulative table score is also based on, at least in part, a cost of the database queries referencing said at least one table.

Claim 8 (previously presented): The method of claim 1 wherein the query statistics comprise column information, and when a column specified in the column information is part of an index, generating an index portion of at least one statistics collection task to collect database statistics for that column.

Claim 9 (previously presented): An article of manufacture comprising a computer usable storage medium embodying one or more instructions executable by a computer for performing a method of collecting statistics in a database management system comprising a

plurality of collection objects having a plurality of tables, at least one particular collection object of the plurality of collection objects comprising a plurality of particular tables of said plurality of tables, the method comprising:

receiving a workload comprising a plurality of database queries referencing the plurality of tables, at least two of the database queries that are different referencing a same table of the plurality of tables;

for each database query of the plurality of database queries of the workload:

identifying query statistics for said each database query of the plurality of database queries, the query statistics comprising a number of references to at least one table of said each database query and a number of occurrences of said each database query; and

consolidating, for said each database query, the query statistics to provide consolidated statistics comprising at least one cumulative table score for said at least one table, respectively, wherein said at least one cumulative table score for said at least one table is based on a previous at least one cumulative table score, if any, for said at least one table and a product of the number of references to said at least one table of said each database query and the number of occurrences of said each database query; and

generating a plurality of statistics collection tasks to collect database statistics to be used by an optimizer based on the consolidated statistics, the plurality of statistics collection tasks to collect the database statistics on the plurality of tables referenced by the plurality of database queries of the workload, the plurality of statistics collection tasks being based on a plurality of collection objects comprising the plurality of tables referenced by the plurality of database queries of the workload, a plurality of collection object scores are determined for the plurality of collection objects comprising the plurality of tables referenced by the plurality of database queries of the workload, each collection object score for each collection object is based on said at least one cumulative table score for said at least one table of said each collection object, wherein the plurality of statistics collection tasks are ordered based on the plurality of collection object scores.

Claim 10 (canceled).

Claim 11 (canceled).

Claim 12 (previously presented): The article of manufacture of claim 9 wherein the collection objects are table spaces.

Claim 13 (canceled).

Claim 14 (canceled).

Claim 15 (previously presented): The article of manufacture of claim 9 wherein the at least one cumulative table score is also based on, at least in part, a cost of the database queries referencing said at least one table.

Claim 16 (previously presented): The article of manufacture of claim 9 wherein the query statistics comprise column information, and when a column specified in the column information is part of an index, generating an index portion of at least one statistics collection task to collect database statistics for that column.

Claim 17 (previously presented): An apparatus for collecting statistics in a database management system comprising a plurality of collection objects having a plurality of tables, at least one particular collection object of the plurality of collection objects comprising a plurality of particular tables of the plurality of tables, comprising:

- a processor; and

- a memory storing one or more instructions that:

- receive a workload comprising a plurality of database queries referencing the plurality of tables, at least two of the database queries that are different referencing a same table of the plurality of tables;

- for each database query of the plurality of database queries of the workload:

identify query statistics for said each database query of the plurality of database queries, the query statistics comprising a number of references to at least one table of said each database query and a number of occurrences of said each database query; and

consolidate, for said each database query, the query statistics to provide consolidated statistics comprising at least one cumulative table score for said at least one table, respectively, wherein said at least one cumulative table score for said at least one table is based on a previous at least one cumulative table score, if any, for said at least one table and a product of the number of references to said at least one table of said each database query and the number of occurrences of said each database query; and

generate a plurality of statistics collection tasks to collect database statistics to be used by an optimizer based on the consolidated statistics, the plurality of statistics collection tasks to collect the database statistics on the plurality of tables referenced by the plurality of database queries of the workload, the plurality of statistics collection tasks being based on a plurality of collection objects comprising the plurality of tables referenced by the plurality of database queries of the workload, a plurality of collection object scores are determined for the plurality of collection objects comprising the plurality of tables referenced by the plurality of database queries of the workload, each collection object score for each collection object is based on said at least one cumulative table score for said at least one table of said each collection object, wherein the plurality of statistics collection tasks are ordered based on the plurality of collection object scores.

Claim 18 (canceled).

Claim 19 (canceled).

Claim 20 (previously presented): The apparatus of claim 17 wherein the collection objects are tablespaces.

Claim 21 (original): The method of claim 1 wherein the at least one cumulative table score represents a total number of references to said at least one table by the plurality of database queries of the workload.

Claim 22 (original): The method of claim 1 wherein said query statistics further comprise a number of references to at least one index of referenced by said each database query of the workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative index score for the at least one index referenced by said each database query, respectively, the at least one cumulative index score being based on a previous at least one cumulative index score, if any, for said at least one index and a product of the number of references to said at least one index and the number of occurrences of said each database query;

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any index referenced by said each database query based on the cumulative index score for that index.

Claim 23 (original): The method of claim 1 wherein said query statistics further comprise a number of references to at least one column-group of said each query of said workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative column-group score for the at least one column-group referenced by said each database query, respectively, said at least one cumulative column-group score being based on a previous at least one cumulative column-group score, if any, for said at least one column-group and a product of the number of references to said at least one column-group and the number of occurrences of said each database query; and

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any table comprising any

column-group referenced by said each database query based on the cumulative column-group score for that column-group.

Claim 24 (currently amended): The method of claim 1 wherein said query statistics further comprise a number of references to at least one column of said each query of said workload;

7L
OK to enter
wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative ~~column-group~~ column score for the at least one column referenced by said each database query, respectively, said at least one cumulative column score being based on a previous at least one cumulative column score, if any, for said at least one column and a product of the number of references to said at least one column and the number of occurrences of said each database query; and

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any table comprising any column referenced by said each database query based on the cumulative column score for that column.

Claim 25 (original): The article of manufacture of claim 9 wherein the at least one cumulative table score represents a total number of references to said at least one table by the plurality of database queries of the workload.

Claim 26 (original): The article of manufacture of claim 9 wherein said query statistics further comprise a number of references to at least one index of referenced by said each database query of the workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative index score for the at least one index referenced by said each database query, respectively, the at least one cumulative index score being based on a previous at least one cumulative index

score, if any, for said at least one index and a product of the number of references to said at least one index and the number of occurrences of said each database query;

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any index referenced by said each database query based on the cumulative index score for that index.

Claim 27 (original): The article of manufacture of claim 9 wherein said query statistics further comprise a number of references to at least one column-group of said each query of said workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative column-group score for the at least one column-group referenced by said each database query, respectively, said at least one cumulative column-group score being based on a previous at least one cumulative column-group score, if any, for said at least one column-group and a product of the number of references to said at least one column-group and the number of occurrences of said each database query; and

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any table comprising any column-group referenced by said each database query based on the cumulative column-group score for that column-group.

Claim 28 (currently amended): The article of manufacture of claim 9 wherein said query statistics further comprise a number of references to at least one column of said each query of said workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative column ~~column-group~~ score for the at least one column referenced by said each database query, respectively, said at least one cumulative column score being based on a previous at least one cumulative column score, if any, for said at least one column and a product of the

TL
ok to
enter

number of references to said at least one column and the number of occurrences of said each database query; and

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any table comprising any column referenced by said each database query based on the cumulative column score for that column.

Claim 29 (original): The apparatus of claim 17 wherein said query statistics further comprise a number of references to at least one index of referenced by said each database query of the workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative index score for the at least one index referenced by said each database query, respectively, the at least one cumulative index score being based on a previous at least one cumulative index score, if any, for said at least one index and a product of the number of references to said at least one index and the number of occurrences of said each database query;

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any index referenced by said each database query based on the cumulative index score for that index.

Claim 30 (currently amended): The apparatus of claim 17 wherein said query statistics further comprise a number of references to at least one column of said each query of said workload;

wherein said consolidating, for said each database query, the query statistics to provide the consolidated statistics comprises determining at least one cumulative column ~~column-group~~ score for the at least one column referenced by said each database query, respectively, said at least one cumulative column score being based on a previous at least one cumulative column score, if any, for said at least one column and a product of the

TL
ok to
enter

Docket No.: SVL920030141US1

Appl. No.: 10/824,856

Amdt. Under 37 CFR § 1.312 dated October 15, 2007

number of references to said at least one column and the number of occurrences of said each database query; and

wherein said generating said at least one statistics collection task also determines the collection object score for the collection object that contains any table comprising any column referenced by said each database query based on the cumulative column score for that column.